

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 777 363 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
04.06.1997 Bulletin 1997/23

(51) Int Cl. 6: H04L 12/56

(21) Application number: 96308420.7

(22) Date of filing: 21.11.1996

(84) Designated Contracting States:
DE FR GB

• Donnelly, Jeffrey M.
Holmdel, NJ 07733 (US)

(30) Priority: 28.11.1995 US 563477

(71) Applicant: NCR International, Inc.
Dayton, Ohio 45479 (US)

(74) Representative: Robinson, Robert George
International Intellectual Property Department,
NCR Limited,
206 Marylebone Road
London NW1 6LY (GB)

(72) Inventors:

• Chien, Anthony H.
Holmdel, NJ 07733 (US)

(54) Method for acknowledgement-based flow control

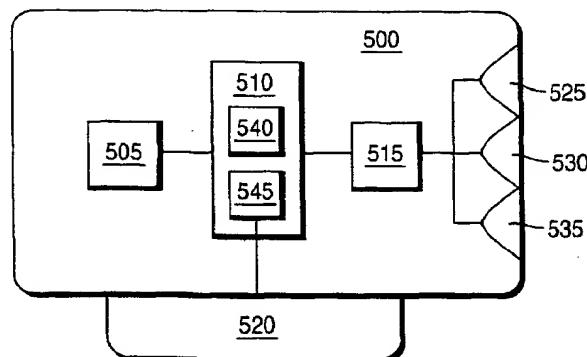
(57) A method and system for intelligent acknowledgement-based flow control in a processing system network. A management circuit (510) governs transmission of data packets (205) and reception indicia by a transmission circuit (515) over a network (100). Detector circuitry (540) is included and is operative to:

- monitor a first latency characteristic of the network (100) being indicative, at least in part, of a utilization level of the network (100); and
- monitor a second latency characteristic being indicative, at least in part, of an efficiency level associated with transmission of the reception indicia by the transmission circuit (515).

Control circuitry (545) is also included and is operative to:

- adjust a retransmission delay of data packets (205) over the network (100) as a function of the first latency characteristic, thereby allowing the management circuit (510) to manage the retransmission delay as a function of the utilization level of said network (100); and
- adjust a transmission delay of reception indicia over the network (100) as a function of the second latency characteristic to thereby allow the management circuit (510) to manage the transmission delay as a function of the efficiency level associated with the transmission circuit (515).

FIG. 5



Description

The present invention relates to network connectivity, and more particularly to circuits and methods for intelligent acknowledgment-based flow control in processing system networks.

A processing system network, such as a computer network or a communications network, is a combination of two or more independent nodes that are capable of communicating with one another over a communications channel, path or link. A node may be an independent processing system, such as a conventional computer, or another processing system network.

Nodes communicate with one another in order to share resources, such as databases and data files, software applications, hardware peripherals and communication links. Communication links are shared to enable two or more nodes that are not directly linked together to communicate through one or more intermediate nodes. Resource sharing generally involves the transference of large amounts of data. The data is typically divided into packets, frames, groups, etc. ("data packets"). Each data packet includes data and information necessary to route the same between two or more nodes.

Networks traditionally belong in one of two general categories, namely, local area networks ("LANs") and wide area networks ("WANs"). A LAN is a group of communicating nodes that are located relatively close to one another, such as within the same building or building complex. A WAN, on the other hand, is a collection of independent and distinct network nodes that cooperate over relatively long distances. Communication links between WAN nodes are routinely provided by third-party carriers, such as by long-distance telephone companies. Gateways, such as routers, bridges or other suitable network portal devices, are used to couple LANs and WANs together (i.e., LAN-to-LAN, LAN-to-WAN and WAN-to-WAN). The conventional approach is to use the gateway as a junction point through which data packets received from a source network are routed to one or more destination networks. Gateways typically include control circuitry and a memory. The memory often includes the routing address for each LAN and each WAN coupled to the gateway, and also often includes the address of one or more of the nodes of each of the LANs and WANs. In the event that a particular node is itself a LAN or a WAN, the addresses of one or more of its nodes are also frequently stored in the memory.

Communication between networks and nodes involves level operations. Level operations include data traffic flow control, sequencing, transmission error detection and correction, and the like. For example, consider three serially coupled networks (i.e., Network A, Network B and Network C) in which one or more data packets are transmitted from a Network A node to a Network C node. By definition, the data packets pass, like a baton in a relay race, from Network A to Network B to

Network C until the same reaches a destination node. The destination node processes the received packets to determine whether they were correctly received. In response thereto, the destination node typically returns either an "ACK" signal, meaning that the transmission was received without error, or a "NAK" signal, meaning that the transmission was corrupt. In the event that the transmitting node receives an ACK signal, it typically purges the previously transmitted data packets from its transmission queue. Alternatively, if it receives a NAK signal, the transmitting node retransmits at least one of the one or more data packets to the destination node.

Timers are typically used to verify the occurrence of an event within a prescribed time period, such as the reception of one or more data packets. When the transmission node transmits a data packet, it starts a transmission timer that is set to expire if the destination node fails to respond in the form of an ACK or a NAK signal within a set period. Upon expiration, the transmission is reattempted and the transmission timer is reset. A common problem with transmission timer values however is that they fail to account for increases and decreases in network traffic (i.e., channel utilization). As traffic increases, timeouts tend to occur too soon. Conversely, as traffic decreases, timeouts tend to occur too late. A further problem associated with destination nodes, is that after returning an ACK, they are often required to return one or more data packets to the transmission node. In other words, after processing the received data packets, the destination node transmits one or more resulting data packets to the transmission node thereby increasing overall network traffic.

A solution to this latter identified problem is to delay return of a positive acknowledgment for a sufficient period for the destination node to process the received data packets, and to "piggyback" the ACK on a data packet to be returned. Typically, when the destination node receives a data packet, it starts an acknowledgment timer that is set to expire if the destination node fails to generate a return data packet within a set period. Upon expiration, the destination node transmits the ACK signal. A common problem with acknowledgment timers is that they fail to account for series of data packets that are received at the destination node that do not generate return data. In other words, the acknowledgment timer tends to timeout repeatedly, unnecessarily delaying transmission of the ACK signal to the transmission node. Unfortunately, these repeated timeouts, when coupled with moderate to heavy traffic over a network, will slow down the transmission mode and may increase the occurrence of transmission timeouts and unnecessary retransmissions.

There accordingly exists a need in the art to effectively limit the occurrence of acknowledgment-based timeouts, particularly, those timeouts resulting from increased network traffic. There exists a further need in the art to more efficiently retransmit lost data packet transmissions that occur during intervals of decreased

network traffic. There exists a still further need in the art to identify series of received data packets that will likely not return data to a transmitting node, and to suitably modify acknowledgment timers to more efficiently acknowledge data packet reception.

It is an object of the present invention to dynamically and adaptively improve the efficiency of data packet transference in order to address the above discussed deficiencies.

According to one aspect of the present invention there is provided a method for governing transmission of a signal by a transmission circuit over a network, characterized by the steps of monitoring a latency characteristic, said latency characteristic being indicative, at least in part, of either an efficiency level associated with transmission of said signal by said transmission circuit, or a utilization level of the network; and adjusting a transmission delay of said transmission circuit for said signal over the network as a function of the latency characteristic to thereby manage the transmission delay.

According to another aspect of the present invention there is provided a management circuit for governing transmission of data packets and reception indicia by a transmission circuit over a network, characterized by detector circuitry operative to monitor a first latency characteristic of the network, said latency characteristic being indicative, at least in part, of a utilization level of the network, and to monitor a second latency characteristic being indicative, at least in part, of an efficiency level associated with transmission of the reception indicia by said transmission circuit; and control circuitry, associated with said detector circuitry and said transmission circuit, operative to adjust a retransmission delay of said transmission circuit of said data packets over the network as a function of the first latency characteristic to thereby allow said management circuit to manage the retransmission delay as a function of said utilization level of said network, and adjust a transmission delay of said transmission circuit of the reception indicia over the network as a function of said second latency characteristic to thereby allow said management circuit to manage said transmission delay as a function of the efficiency level associated with said transmission circuit.

One embodiment of the present invention will now be described by way of example with reference to the accompanying drawings in which:-

Fig. 1 is a block diagram of the architectural structure of an exemplary transport control protocol/internet protocol (TCP/IP) network;

Fig. 2 is a block diagram for routing one or more data packets between first and second locations in another exemplary processing system network;

Fig. 3 is an isometric view of an exemplary processing system capable of functioning as a node within a processing system network;

Fig. 4 is a block diagram of an exemplary micro-processing system that may suitably be implement-

ed in the processing system of Fig. 3;

Fig. 5 is a high-level block diagram of an exemplary network interface in accordance with the principles of the present invention; and

Fig. 6 is a flow diagram for performing intelligent acknowledgment-based flow control in accordance with the principles of the present invention.

Referring now to Fig. 1, a block diagram of the architectural structure of an exemplary TCP/IP network 100 is illustrated. TCP/IP is a known set of protocols developed to link dissimilar computers across many kinds of networks. TCP/IP is used herein for illustrative purposes only. In point of fact, the principles of the present invention may be implemented in accordance with any suitable acknowledgment-based data packet transmission scheme, including open systems interconnection (OSI), X.25, and the like. "Include," as the term is used herein, means "include without limitation."

Exemplary network 100 includes a plurality of interconnected nodes 105. A first set of nodes 105 form a first sub-network 110a, while a second set of nodes 105 form a second sub-network 110b. Exemplary sub-networks 110 are coupled via an exemplary gateway 115 (i.e., router, bridge or other suitable portal device). "Or" as used herein means "and/or." Gateway 115 may suitably be one of a core gateway or a non-core gateway. A core gateway includes information about the structure of network 100, whereas a non-core gateway includes limited or incomplete routing information, i.e., knowledge regarding one or more routes to a core gateway, but no knowledge about routing beyond the same.

Core gateways typically store routing information in one or more routing tables. Each node 105 is required only to know a specific route to its local gateway 115. A data packet transmitted between two remotely located nodes is accordingly passed from one gateway to another until the same reaches a destination node.

A routing table may suitably be programmer defined or developed dynamically (i.e., gateway queries neighboring gateways or the like for routing information). A given gateway may suitably include routing information for directly attached networks and general knowledge for routing traffic to remote networks. If a gateway cannot resolve a particular routing address, it broadcasts a request seeking assistance concerning the same to other gateways.

In general, as a data packet is forwarded between locations spanning a plurality of sub-network environments, it is processed to determine the address of its destination location or, possibly, intermediate locations. In response thereto, the data packet is suitably encased or wrapped in a data envelope. The data envelope, such as a file transfer protocol frame, typically includes a header having routing and transmission information. A TCP header, for example, provides end-to-end transmission confirmation or acknowledgment.

Referring next to Fig. 2, a block diagram for routing

one or more data packets between first and second locations in an exemplary TCP/IP network 100 is illustrated in accordance with the present invention. Exemplary network 100 includes a plurality of nodes 105, each of which is coupled locally to a suitable sub-network 110. A first node 105a, operating as an internet protocol (IP) host, transmits an exemplary data packet 205 to a destination node 105n using an IP addressing scheme.

A first sub-network 110a processes data packet 205 and determines that destination node 105n is not local. In response thereto, an IP header is added to data packet 205. The IP header includes a global IP address for a second sub-network 110n. Gateway 115 routes data packet 205 based on the address of second sub-network 110n. Second sub-network 110n determines if the address of destination node 105n is local or remote. If the same is local, as illustrated, the data envelope is stripped off and data packet 205 is routed to destination node 105n.

In the event that destination node 105n was not local, network 110n would suitably determine the address of the next network element and add its address to the existing IP data packet. When that element received the data envelope containing the IP data packet, it would strip off the header designating its own address, and go through the same process as the previous network did, determining whether the destination address was local, and if not, packaging the IP data packet within another addressing header for the next network element. This process would continue until the data packet was finally routed to the destination address.

In the illustrated embodiment, first node 105a is further operative to retransmit data packet 205 in response to a transmission timeout. A "timeout," as the term is used herein, includes any suitable comparison between a threshold and a timer or clock value that yields a particular result. Destination node 105n is operative in accordance with TCP to verify and, preferably, correct received data packet 205, if necessary and possible. If correct, destination node 105n returns an acknowledgment signal to first node 105a, otherwise it may suitably return a negative acknowledgment.

Destination node 105n is also preferably operative to suitably use an acknowledgment timer to measure an interval of elapsed time with respect to the reception of data packet 205. "Measure," as used herein, includes compare, count, equate, evaluate, gauge, survey, quantify, weigh or the like. In the event that the processing data packet 205 generates or results in one or more data packets to be returned to first node 105a, destination node piggy-backs an acknowledgment, such as an ACK, to at least one return data packet. Destination node 105n is preferably further operative to transmit the acknowledgment alone to first node 105a in response to an acknowledgment timeout.

Exemplary transmission timer values may suitably depend, without limitation, upon one or more of the following:

(a) round trip propagation delay of the signal (usually a small value, except for very long and very high-speed circuits);

5 (b) processing time at the receiver (including queuing time of the data packet);

(c) transmission time of the acknowledging signal or frame; or

(d) possible queue and processing time at the transmitter when it receives the acknowledgment.

10 Recall that heavy network traffic may delay receipt of the acknowledgment signal by first node 105b, thereby causing a timeout and retransmission of data packet 205. First node 105a is operative in accordance with the 15 illustrated embodiment to monitor the frequency of transmission timeout occurrences. If the same compares unfavourably with a threshold, then first node 105a is preferably operative to modify the timeout function (e.g., change the timeout threshold, slow the timer or clock, etc.); temporarily disable the timeout function, temporarily disable the acknowledgment requirement, or the like. "Disable," as the term is used herein, includes detach, disassociate, disconnect, disengage, impair, incapacitate, separate or the like. An important aspect of 20 the illustrated embodiment is the improved efficiency of data packet transference.

25 Destination node 105n is operative in accordance with the illustrated embodiment to monitor the frequency of acknowledgment timeout occurrences. If the same compares unfavourably with a threshold, then destination node 105n is preferably operative to modify the timeout function (e.g., change the timeout threshold, slow the timer or clock, etc.), temporarily disable the timeout function, temporarily disable the acknowledgment requirement, or the like. Another important aspect of 30 the illustrated embodiment is the improved efficiency of acknowledgment transference.

35 Referring to Fig. 3, an isometric view of an exemplary processing system 105 is illustrated. Processing system 105 is capable of functioning as a node within either of the exemplary processing system networks of Figs. 1 or 2. Processing system 105 includes a monitor 305, a housing 310 and a keyboard 315.

40 Housing 310 includes a hard disk drive 320 and a floppy disk drive 325. Hard disk drive 320 is suitable to provide fast access storage and retrieval. Floppy disk drive 325 is operative to receive, read or write to external disks, and may suitably be replaced by or combined with other conventional structures for transferring data or instructions, including tape and compact disc drives, telephony systems and devices (including telephone, video phone, facsimile or the like), message paging, network communication ports and the like.

45 Housing 310 is partially cut-away to illustrate a battery 330, a clock 335, a processor 340 and a detached local memory 345, all of which are suitably housed therein. Although processing system 105 is illustrated having a single processor, a single hard disk drive and

a single local memory, processing system 105 may suitably be equipped with any multitude or combination of processors or storage devices. Processing system 105 may, in point of fact, be replaced by, or combined with, any suitable processing system operative in accordance with the principles of the present invention, including video phones, telephones, televisions, pagers, sophisticated calculators, and hand-held, laptop/notebook, mini, mainframe and super computers, as well as processing system network combinations of the same.

Conventional processing system architecture is more fully discussed in Computer Organization and Architecture, by William Stallings, MacMillan Publishing Co. (3rd ed. 1993); conventional processing system network design is more fully discussed in Data Network Design, by Darren L. Spohn, McGraw-Hill, Inc. (1993); and conventional data communications is more fully discussed in Data Communications Principles, by R. D. Gitlin, J. F. Hayes and S. B. Weinstein, Plenum Press (1992) and in The Irwin Handbook of Telecommunications, by James Harry Green, Irwin Professional Publishing (2nd ed. 1992). Each of the foregoing publications is incorporated herein by reference.

Referring to Fig. 4, a block diagram of an exemplary microprocessing system 400 is illustrated that may suitably be implemented in processing system 105 of Fig. 3. Microprocessing system 400 includes a processor 340 coupled via data bus 405 with a detached local memory 345. Memory 345 is operative to store data or instructions, which processor 340 is operable to retrieve and execute.

Processor 340 includes a control unit 410, an arithmetic and logic unit (ALU) 415, and a internal memory 420 (e.g., stackable cache, a plurality of registers, etc.). Control unit 410 is suitably operative to fetch ones of the instructions from memory 345. ALU 415 is suitably operative to perform a plurality of operations, such as addition and Boolean AND, needed to carry out those instructions. Internal memory 420 is suitably operative to provide local high speed storage used to store temporary results and control information.

Referring to Fig. 5, a high-level block diagram of an exemplary network interface 500 (e.g., a modem card, a FDDI card, a Ethernet card, etc.) is illustrated in accordance with the principles of the present invention. Interface 500 includes an illustrative memory 505, a management circuit 510, a transmission circuit 515, an illustrative expansion bus connector 520, an illustrative connector for twisted-pair wire 525, an illustrative coaxial cable connector 530, and an illustrative transceiver connector 535. Management circuit 510 includes detector circuitry 540 and control circuitry 545.

Interface 500 may suitably be coupled with exemplary processing system 105 of Fig. 3. Management circuit 510 is operative to govern transmission of one or more of a plurality of data packets, as well as reception indicia, such as ACKs and NAKs, by transmission circuit 515 over a network, such as network 100 of Figs. 1 and

2. The management circuit 510 comprises detector circuitry 540 and control circuitry 545.

Detector circuitry 540 is operative to monitor a first latency characteristic of the network. The first latency characteristic is indicative or representative of, at least in part, a utilization level of the network. The utilization level of the network is important with respect to data packet transmission, in general, and to data packet retransmission, more specifically. As has been discussed herein, the threshold associated with a transmission timer is critical to efficient retransmission. Detector circuitry 540 is further operative to monitor a second latency characteristic. The second latency characteristic is indicative of, at least in part, an efficiency level associated with transmission of the reception indicia by transmission circuit 515. The efficiency level associated with the transmission of the reception indicia is important with respect to data packet acknowledgment, in general, and to positive acknowledgment and ACK piggybacking, more specifically. As discussed above, the threshold associated with an acknowledgment timer is critical to efficient data packet acknowledgment.

Control circuitry 545 is associated with both detector circuitry 540 and transmission circuit 515. Control circuitry 545 is operative to suitably adjust a retransmission delay associated with transmission of the data packets over the network. The retransmission delay is the time interval between transmission and subsequent retransmission of one or more data packets. The adjustment is performed as a function of the first latency characteristic. This suitably enables or allows management circuit 510 to manage the retransmission delay as a function of the utilization level of the network.

Control circuitry 545 is further operative to suitably adjust a transmission delay associated with transmission of the reception indica over the network. The transmission delay is the time interval between reception of a data packet and the eventual transmission of the reception indicia without piggybacking a return data packet. The adjustment is performed as a function of the second latency characteristic. This suitably enables management circuit 510 to manage the transmission delay as a function of the efficiency level associated with transmission circuit 515.

In performing either of the two above-identified adjustments, control circuitry 545 may suitably use principles of applied mathematical theories, including statistics, stochastic modelling, chaos theory, standard deviation, probability theory, permutations and combinations, frequency, or the like. Further, in deriving the first measurable characteristic, in addition to the use of timers, clocks and the like, any physically sensible characteristic or aspect of the network impacting communication throughput may suitably be used (e.g., communication channel utilization indicia).

Although network interface 500 is used to illustrate one circuit embodiment of the present invention, other circuit configurations may suitably be implemented in

other processing systems, nodes, gateways or the like. More particularly, in alternate embodiments, the above-identified circuits and circuitry, as well as microprocessing system 400 of Fig. 4, may suitably be replaced by or combined with programmable logic devices, such as PALs (programmable array logic) and PLAs (programmable logic arrays), DSPs (digital signal processors), FPGAs (field programmable gate arrays), ASICs (application specific integrated circuits), VLSIs (very large scale integrated circuits) or the like.

Referring to Fig. 6, a flow diagram is illustrated for performing intelligent acknowledgment-based flow control in accordance with the principles of the present invention. For illustrative purposes only, the discussion of Fig. 6 is made with reference to Fig. 5. However, the management circuit 510 and the transmission circuit 515 could also be located in nodes.

A determination is made as to whether a data packet has been received from a transmission node, gateway or the like (decisional step 602). In the event a data packet was received (YES branch of decisional step 602), the transmission circuit 515 is disabled, an acknowledgment timer is reset and the acknowledgment timer begins generating a transmission delay (process step 604). The transmission delay represents an elapsed time interval with respect to a received data packet. A determination is made as to whether the transmission delay compares unfavourably with a transmission threshold (decision step 606).

In the event that the transmission delay compares unfavourably (YES branch of decisional step 606), then a reception indicia, such as a signal, is returned to the transmission node, gateway or the like (input/output step 608). Detector circuitry 540 is operative to derive the latency characteristic as a function of the comparison (process step 610), and to compare the same with a latency threshold (process step 612), such as a maximum frequency of occurrence, for example. In the event that the latency characteristic compares unfavourably (YES branch of decisional step 614), then control circuitry 545 is operative to adjust or modify the acknowledgment timeout function (e.g., change the timeout threshold, slow the timer or clock, etc.), temporarily disable the same, temporarily disable the acknowledgment requirement, or the like (process step 616), thereby improving the overall efficiency of acknowledgment transference.

In other words, management circuit 510 is operative to monitor the latency characteristic, and to adjust the transmission delay of the transmission of the reception indicia over the network. This is preferably accomplished as a function of the latency characteristic.

In the event a data packet was not received (NO branch of decisional step 602), then a determination is made as to whether a data packet is to be transmitted (decisional step 618). In the event the data packet is to be transmitted (YES branch of decisional step 618), then the transmission circuit 515 is enabled, a retrans-

mission timer is reset and the retransmission timer begins generating a retransmission delay (process step 620). The retransmission delay represents an elapsed time interval with respect to transmission of the data packet without receipt of reception indicia, such as an acknowledgment signal, for example. A determination is made as to whether the retransmission delay compares unfavourably with a retransmission threshold (decision step 622).

- 10 In the event that the retransmission delay compares unfavourably (YES branch of decisional step 622), then the data packet is retransmitted over the network (input/output step 624). Detector circuitry 540 is operative to derive the latency characteristic as a function of the comparison (process step 610), and to compare the same with a latency threshold (process step 612), such as a maximum frequency of occurrence, for example. In the event that the latency characteristic compares unfavourably (YES branch of decisional step 614), then control circuitry 545 is operative to adjust or modify the retransmission timeout function (e.g., change the timeout threshold, slow the timer or clock, etc.), temporarily disable the same, temporarily disable the retransmission requirement, or the like (process step 626), thereby improving the overall efficiency of data packet retransmission.

The principles of the present invention, as illustrated above, may be suitably implemented as software. An exemplary software embodiment includes a plurality of

- 30 instructions stored to a conventional storage medium. The instructions are readable and executable by a suitable processor. The instructions, upon execution, operate to control the processor to route data packets through a processing system network in accordance with the present invention. Preferred storage media include without limitation, magnetic, optic, and semiconductor, as well as suitable combinations thereof. A preferred software embodiment for governing transmission of reception indicia by a transmission circuit forms a portion of NCR's STARGROUP™ LAN Manager for UNIX™ -- OSI Transport Version, which is available from NCR Corporation in Dayton, Ohio.

45 Claims

1. A method for governing transmission of a signal by a transmission circuit (515) over a network (100), characterized by the steps of:
 - 50 monitoring a latency characteristic, said latency characteristic being indicative, at least in part, of either an efficiency level associated with transmission of said signal by said transmission circuit, or a utilization level of the network; and adjusting a transmission delay of said transmission circuit for said signal over the network as a function of the latency characteristic to there-

by manage the transmission delay.

2. A method according to claim 1, wherein said signal is a reception indicium, characterized in that:

said latency characteristic is indicative, at least in part, of an efficiency level associated with transmission of said reception indicia by said transmission circuit (515); and
 transmission delay is managed as a function of the efficiency level associated with said transmission circuit where said transmission delay represents an elapsed interval with respect to a received data packet (205).

3. A method according to claim 1, wherein said signal is a data packet (205), characterized in that:

said latency characteristic is indicative, at least in part, of a utilization level of said network (100); and
 transmission delay is a retransmission delay and is managed as a function of the utilization level of said network where said retransmission delay represents an elapsed interval with respect to a transmitted data packet.

4. A method according to claim 3, characterized in that:

said transmission circuit (515) is disabled in response to a received positive acknowledgment; and
 said transmission circuit is enabled to retransmit said data packets (205) over said network (100) in response to a negative acknowledgment.

5. A method according to any preceding claim, characterized by the step of comparing said transmission delay with a threshold.

6. A method according to claim 5, characterized by the step of deriving said latency characteristic as a function of said comparison.

7. A method according to any preceding claim, characterized by the step of comparing said latency characteristic with a threshold.

8. A management circuit (510) for governing transmission of data packets (205) and reception indicia by a transmission circuit (515) over a network (100), characterized by:

detector circuitry (540) operative to monitor a first latency characteristic of the network, said latency characteristic being indicative, at least in part, of a utilization level of the network, and to monitor a second latency characteristic being indicative, at least in part, of an efficiency level associated with transmission of the reception indicia by said transmission circuit; and control circuitry (545), associated with said detector circuitry and said transmission circuit, operative to adjust a retransmission delay of said transmission circuit of said data packets over the network as a function of the first latency characteristic to thereby allow said management circuit to manage the retransmission delay as a function of said utilization level of said network, and adjust a transmission delay of said transmission circuit of the reception indicia over the network as a function of said second latency characteristic to thereby allow said management circuit to manage said transmission delay as a function of the efficiency level associated with said transmission circuit.

9. A management circuit (510) according to claim 8 characterized in that said management circuit is associated with one of a node (105) and gateway (115).

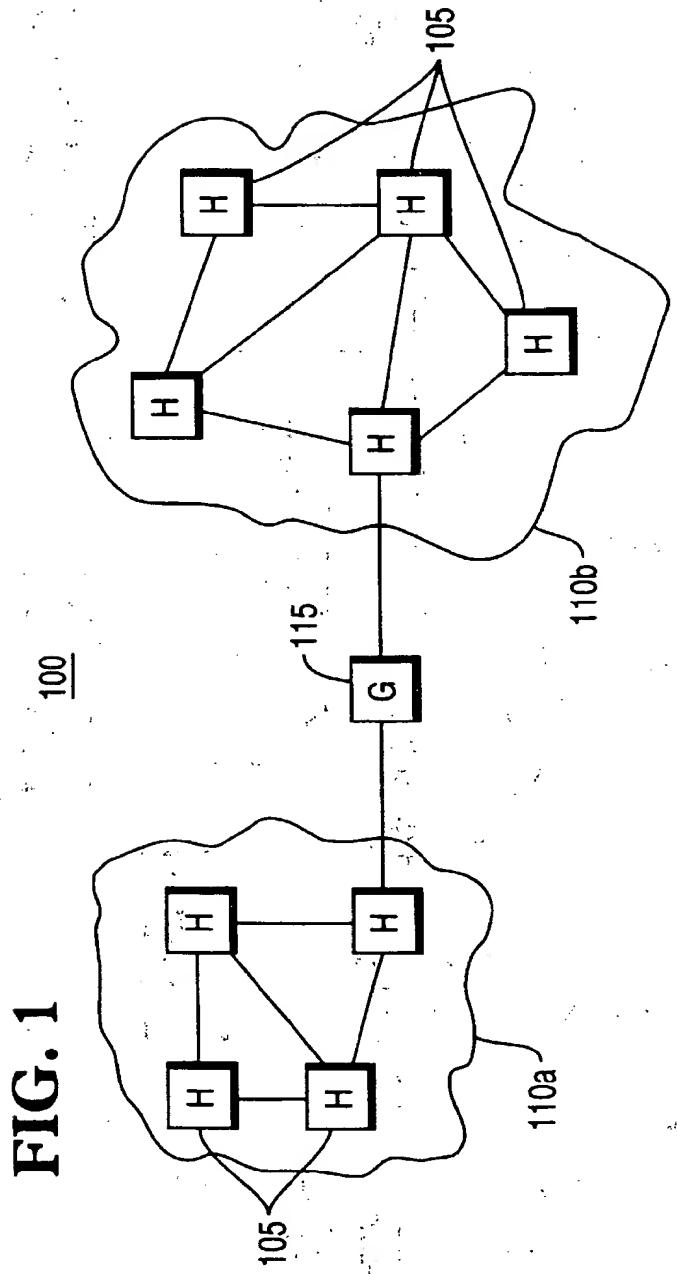


FIG. 1

FIG. 2

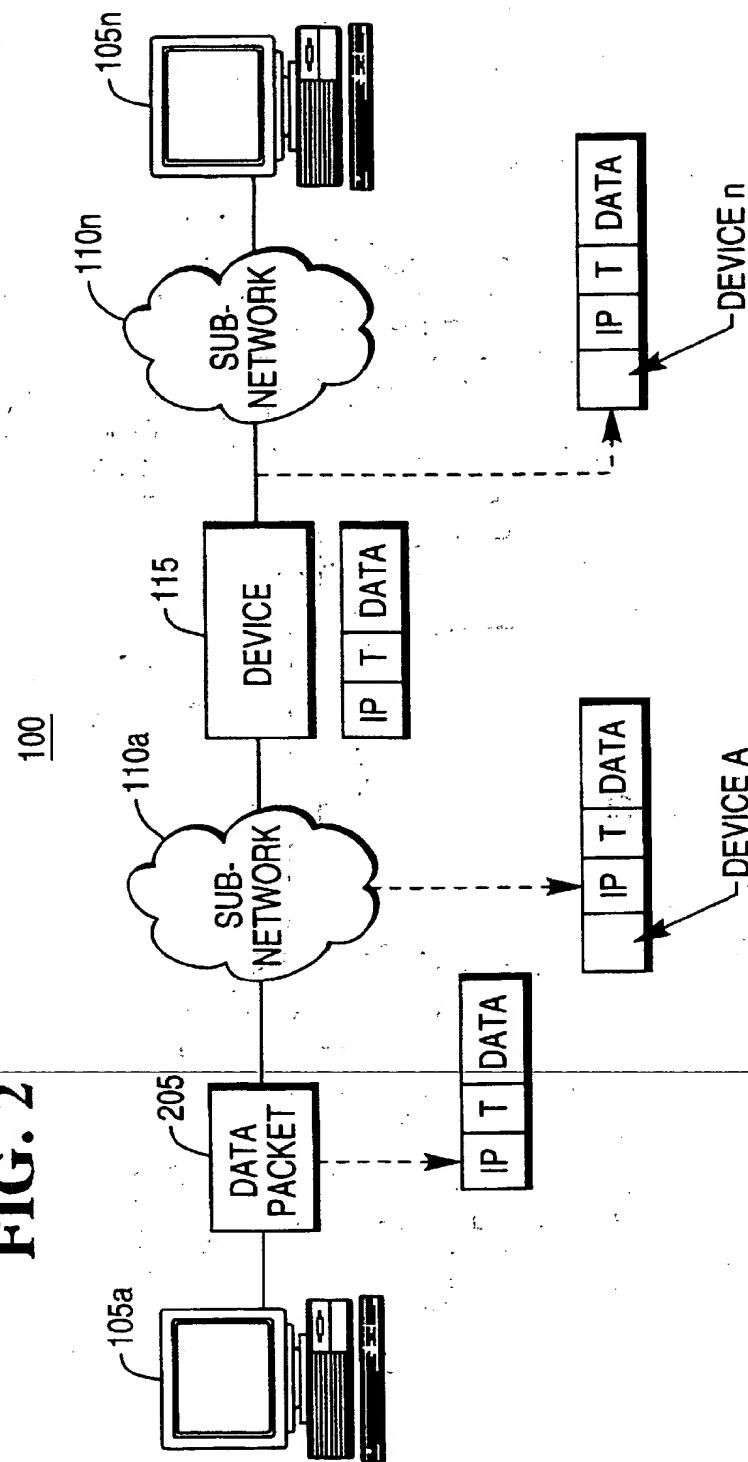


FIG. 3

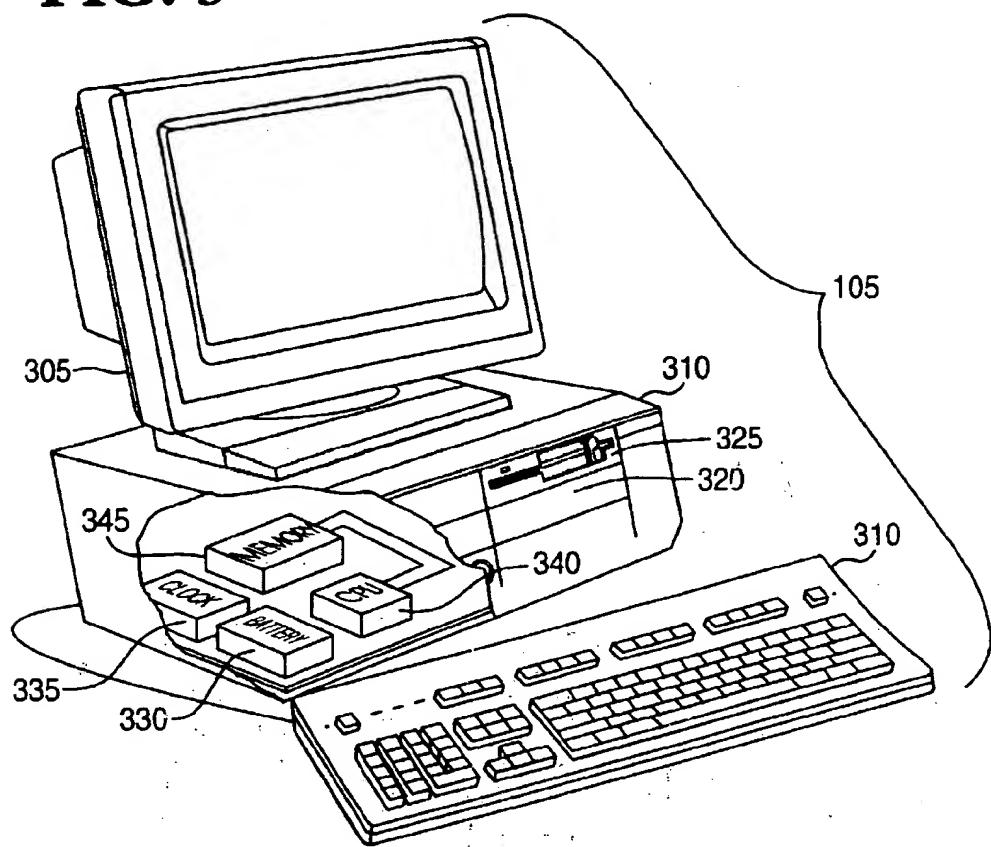


FIG. 4

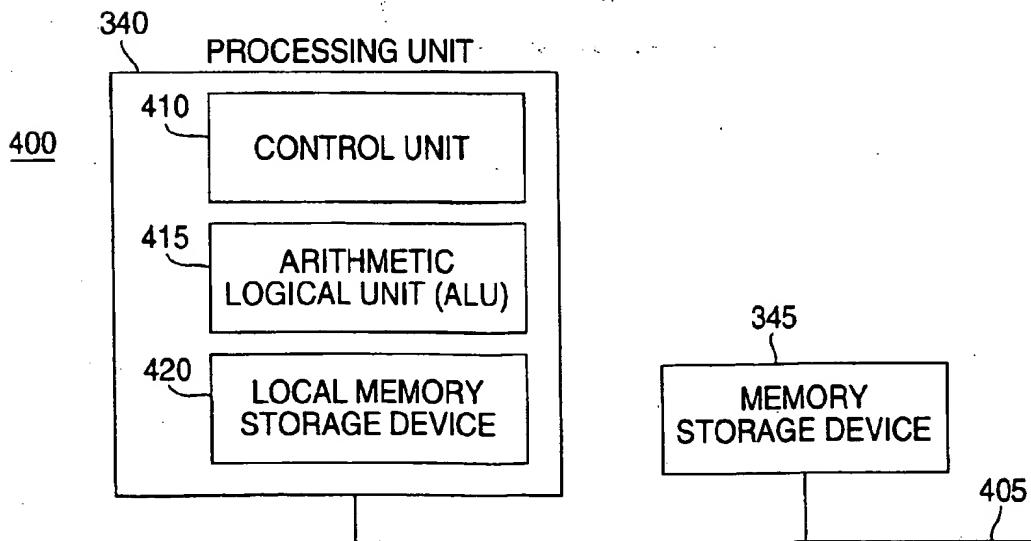


FIG. 5

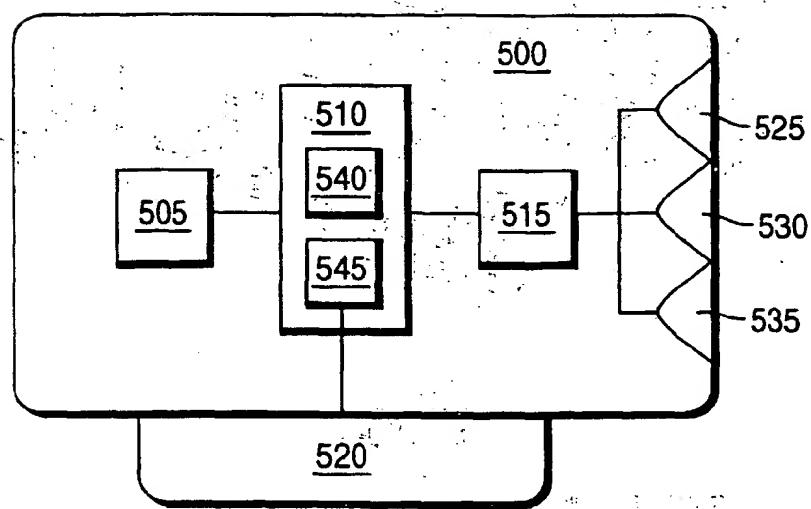
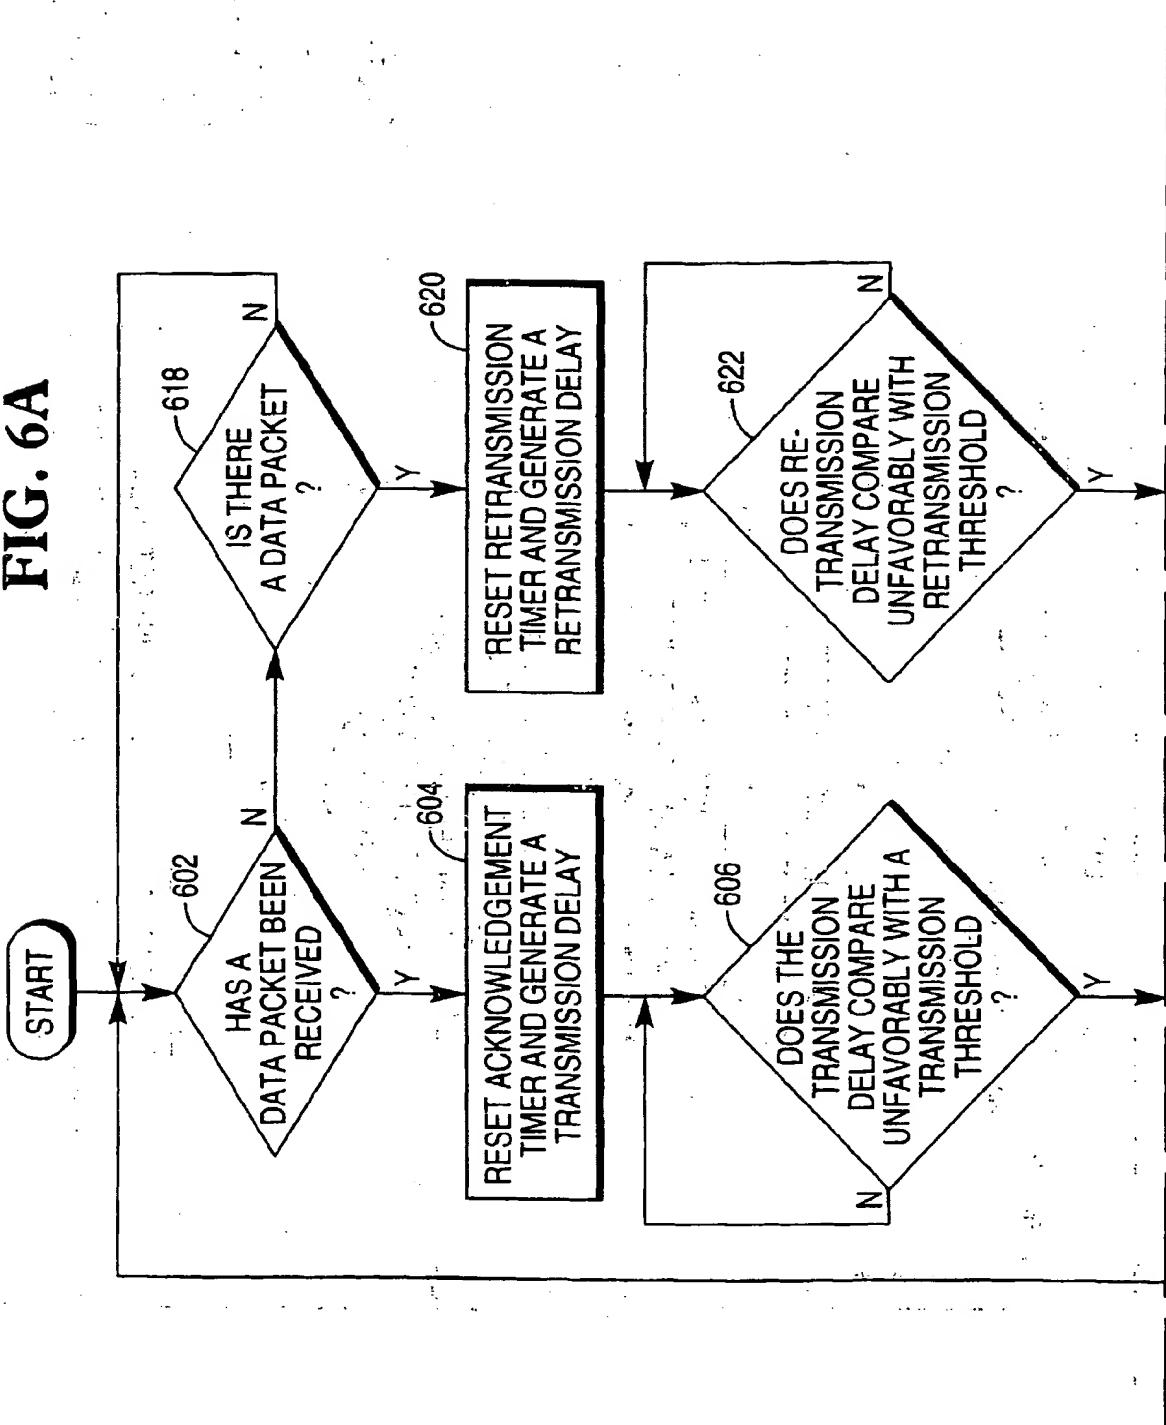
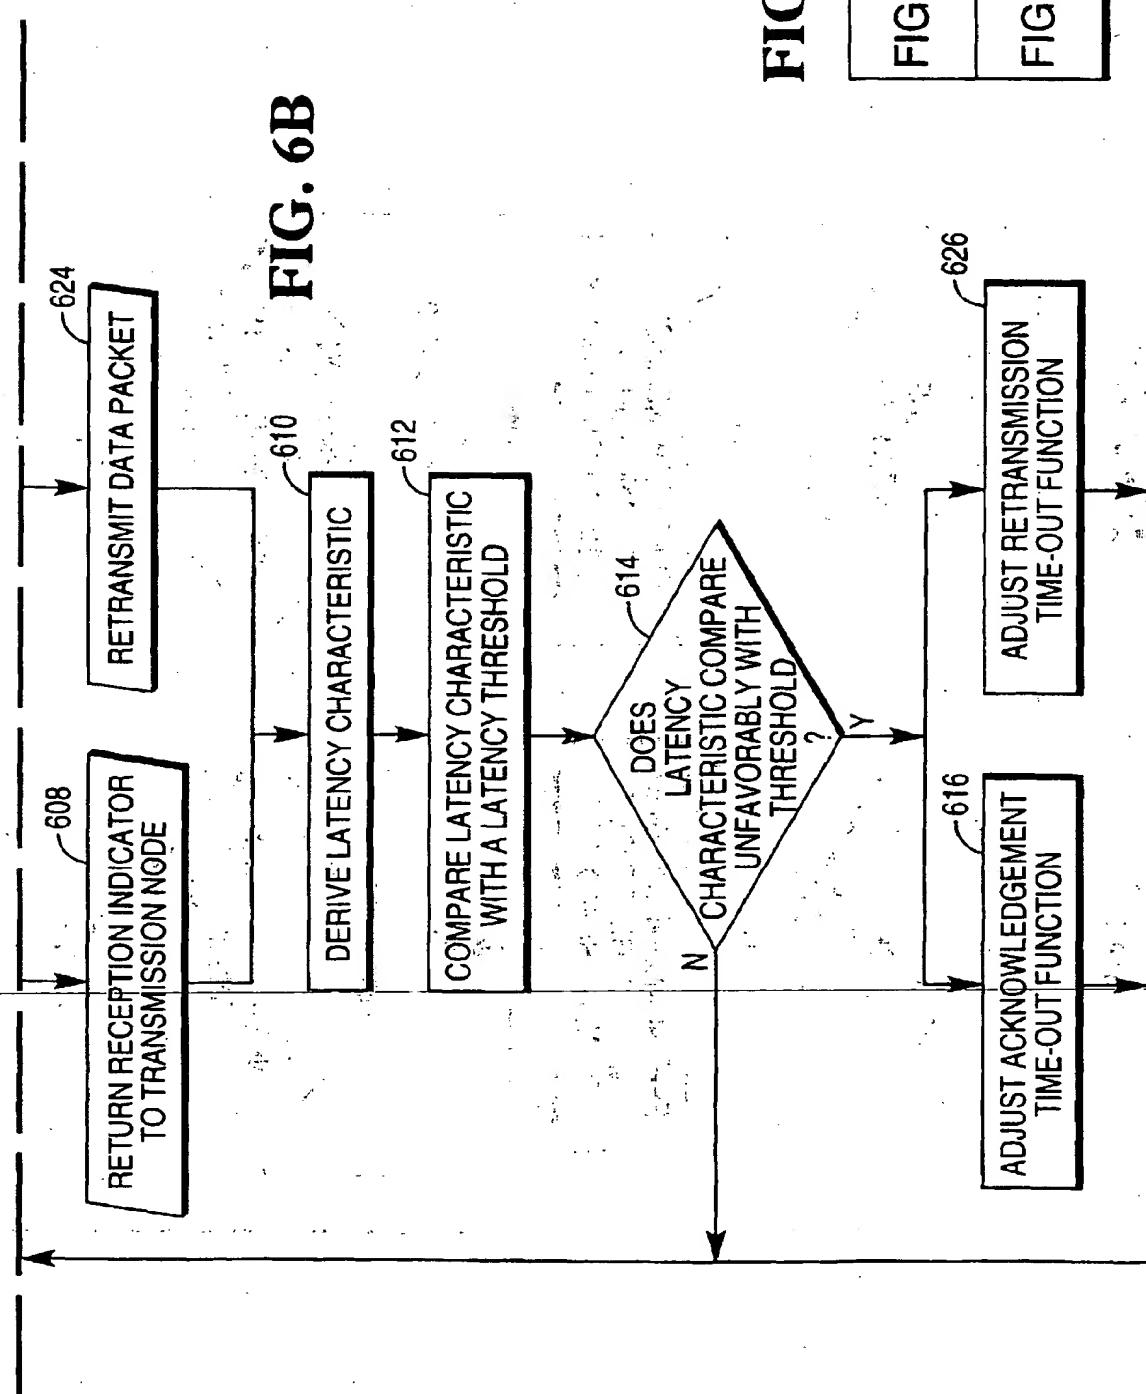
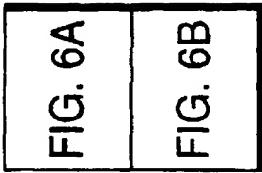


FIG. 6A

**FIG. 6**

THIS PAGE BLANK (USPTO)



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 777 363 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
07.02.2001 Bulletin 2001/06

(51) Int Cl. 7: H04L 12/56

(43) Date of publication A2:
04.06.1997 Bulletin 1997/23

(21) Application number: 96308420.7

(22) Date of filing: 21.11.1996

(84) Designated Contracting States:
DE FR GB

- Donnelly, Jeffrey M.
Holmdel, NJ 07733 (US)

(30) Priority: 28.11.1995 US 563477

(74) Representative: Irish, Vivien Elizabeth
International IP Department,
NCR Limited,
206 Marylebone Road
London NW1 6LY (GB)

(71) Applicant: NCR International, Inc.
Dayton, Ohio 45479 (US)

(72) Inventors:

- Chien, Anthony H.
Holmdel, NJ 07733 (US)

(54) Method for acknowledgement-based flow control

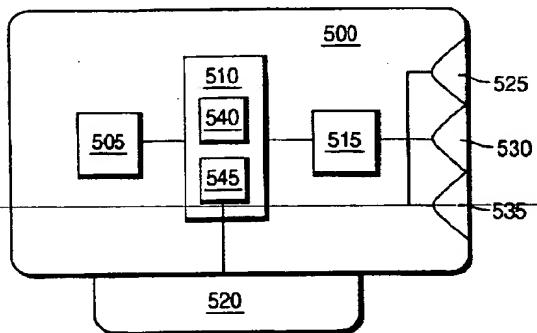
(57) A method and system for intelligent acknowledgement-based flow control in a processing system network. A management circuit (510) governs transmission of data packets (205) and reception indicia by a transmission circuit (515) over a network (100). Detector circuitry (540) is included and is operative to:

- (a) monitor a first latency characteristic of the network (100) being indicative, at least in part, of a utilization level of the network (100); and
- (b) monitor a second latency characteristic being indicative, at least in part, of an efficiency level associated with transmission of the reception indicia by the transmission circuit (515).

Control circuitry (545) is also included and is operative to:

- (a) adjust a retransmission delay of data packets (205) over the network (100) as a function of the first latency characteristic, thereby allowing the management circuit (510) to manage the retransmission delay as a function of the utilization level of said network (100); and
- (b) adjust a transmission delay of reception indicia over the network (100) as a function of the second latency characteristic to thereby allow the management circuit (510) to manage the transmission delay as a function of the efficiency level associated with the transmission circuit (515).

FIG. 5





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 96 30 8420

DOCUMENTS CONSIDERED TO BE RELEVANT									
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)						
X	EP 0 329 159 A (NIPPON ELECTRIC CO) 23 August 1989 (1989-08-23) * column 4, line 17-29 * * column 5, line 3-9 * * column 5, line 22-34 * * column 5, line 37-47 * * column 5, line 55 - column 6, line 4 * * claims 1,2,8,10,11 *	1,2,7 3,4,8	H04L12/56						
X	ZHANG L: "WHY TCP TIMERS DON'T WORK WELL" COMPUTER COMMUNICATIONS REVIEW, US, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, vol. 16, no. 3, 5 August 1986 (1986-08-05), pages 397-405, XPO00743153 ISSN: 0146-4833	1,5,6							
Y	* page 398, paragraph 3 * * page 398, paragraph 3.1. *	3,4,8							
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)						
			H04L						
<p>The present search report has been drawn up for all claims</p> <table border="1"> <tr> <td>Place of search</td> <td>Date of completion of the search</td> <td>Examiner</td> </tr> <tr> <td>THE HAGUE</td> <td>18 December 2000</td> <td>Hardelin, T</td> </tr> </table> <p>CATEGORY OF CITED DOCUMENTS</p> <ul style="list-style-type: none"> X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited for the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>				Place of search	Date of completion of the search	Examiner	THE HAGUE	18 December 2000	Hardelin, T
Place of search	Date of completion of the search	Examiner							
THE HAGUE	18 December 2000	Hardelin, T							

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 96 30 8420

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

18-12-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0329159 A	23-08-1989	JP 2001665 A JP 2540930 B US 5042029 A	05-01-1990 09-10-1996 20-08-1991

THIS PAGE BLANK (USPTO)